Prospects and Limitations of Algorithmic Cooling

Gilles Brassard[1], Yuval Elias[2], Tal Mor[2], and Yossi Weinstein[2]

[1] Département IRO, Université de Montréal, Montréal, H3C 3J7 Canada
[2] Department of Computer Science, Technion 32000, Haifa, Israel

April 29, 2014

arXiv:1404.6824v1 [quant-ph] 27 Apr 2014

Abstract

Heat-bath algorithmic cooling (AC) of spins is a theoretically powerful effective cooling approach, that (ideally) cools spins with low polarization exponentially better than cooling by reversible entropy manipulations alone. Here, we investigate the limitations and prospects of AC. For non-ideal and semioptimal AC, we study the impact of finite relaxation times of reset and computation spins on the achievable effective cooling. We derive, via simulations, the attainable cooling levels for given ratios of relaxation times using two semioptimal practicable algorithms. We expect this analysis to be valuable for the planning of future experiments. For ideal and optimal AC, we make use of lower bounds on the number of required reset steps, based on entropy considerations, to present important consequences of using AC as a tool for improving signal-to-noise ratio in liquid-state magnetic resonance spectroscopy. We discuss the potential use of AC for noninvasive clinical diagnosis and drug monitoring, where it may have significantly lower specific absorption rate (SAR) with respect to currently used methods.

1 Introduction

Algorithmic cooling (AC) [1] is useful for initiating a spin-half system, or more generally, a quantum bit (qubit) or quantum digit (qudit) system [2–7], and for potentially improving signal-to-noise ratio (SNR) in liquid-state magnetic resonance spectroscopy [2, 8–12]. AC of spins allows effective cooling of a single or several spins far beyond close-system cooling (reversible entropy manipulations [13, 14]). Such effective spin-cooling is used to reach an effective temperature, smaller (or even much smaller) than the actual temperature of the environement, and then the spins relax back to the temperature of the environement, in a typical time scale called the relaxation time T_1.

See [2] for a detailed presentation of the bound of $\sqrt{n}$ on cooling n spins using reversible entropy manipulations and the exponential bypass achievable by AC. Exponential cooling was shown for practicable AC (PAC) algorithms [2], and improved exponential cooling was shown for optimal algorithms[1] — the partner-pairing algorithm [3, 4] and all-bonacci [15, 16]. Note that the analysis here is relevant for high-temperature (e.g. room temperature) liquid NMR, hence for a very small initial polarization bias, $\varepsilon \ll 1$ (typically around 10^{-5}). The exponential cooling mentioned above is an approximation which is true as long as the final bias also satisfies that it is much smaller than 1. In principle, AC is a microscopic-scale quantum heat pump [2, 17–23], and it is found useful recently for various purposes, see [24–27], in addition to the abovementioned ones.

AC is based on the presence, within the same spin system, of two types of spins: spins that interact rapidly with the environment, termed *reset spins* and slow-relaxing *computation spins*, which can store and exploit the enhanced polarization. Ideally, the reset spins repolarize infinitely faster than the computation spins. Let $\mathcal{R}$ denote the ratio between the relaxation time of the computation spins (T_1(comp)), and the relaxation time of the reset spins (T_1(reset)), namely, $\mathcal{R} \triangleq T_1(\text{comp})/T_1(\text{reset})$. The ideal scenario where $\mathcal{R} \to \infty$ is assumed in most theoretical papers, e.g., in [1–4, 15, 16, 28]. As suggested in [28], it is important to deal with a realistic scenario where $\mathcal{R}$ is finite, and as suggested in [10] it is interesting to compare AC to other methods used to improve the SNR in magnetic resonance spectroscopy.

The search for better algorithms has lead to both optimal AC [3, 4], and practicable AC (PAC) [2, 4, 15, 16, 28]. The optimal "partner-pairing algorithm" (PPA) [3, 4] was shown to require the least amount of reset steps (namely, WAIT steps), however, its unitary steps are extremely complicated (when written using 2-3 qubit gates). In contrast, PAC [2] and semioptimal PAC (SOPAC) [28], which use very simple quantum gates in-between reset steps, cool various spin-systems to a much lesser extent than the PPA, for any given number of reset steps. For any odd number of spins, $n = 2j+1$, the SOPAC algorithm called mPAC (see [28]) cools, ideally, the most significant bit by a factor of $\left(2 - \frac{1}{2^m}\right)^j$, where m denotes the number of compressions at each recursive level, and j denotes the purification level. For $m = 2$, the exponential cooling of 1.75^j provided by 2PAC is much greater than the exponential cooling [2] provided by 1PAC.

Optimal cooling (via PPA) could approach a cooling factor of 2^{n-2}, and for any non-trivial cooling level (e.g., $n \gtrsim 6$) achievable by both the PPA and 2PAC, the PPA uses much fewer reset steps. Another SOPAC algorithm, named mFib, also presented in [10], belongs to a family of algorithms containing also δ-Fib, Tribonacci and others [4, 15, 16], the bonacci AC, that deserve more attention and are reviewed and extended in section 2.

When AC is moved from theory to experiment [8, 10–12, 29–31] the cooling, of course, is limited due to the realistic reset time. The reset spins do not relax infinitely faster than the computation spins, and one cannot assume anymore that computation spins do not reheat during reset steps. In real life one must account for a finite ratio between the

[1] The word *optimal* here refers to the achieved spin temperature. The partner pairing algorithm is also most efficient with respect to the number of reset steps, see below.

relaxation times of the computation and reset spins. In sections 3 and 4 we analyze two algorithms, the PPA and 2PAC, that ideally provide exponential cooling, in order to present two limitations of AC and understand their consequence.

In section 3 we present numerical simulation of the practicable algorithm 2PAC, applied onto 7 spins, for $100 \leq \mathcal{R} \leq 10,000$. We use a simplified model of relaxation that prevents the need of simulating the relaxation of all correlation terms between spins, and is therefore suitable for a greater number of spins. We show precisely how the exponential advantage over reversible cooling is lost when the number of spins is increased or $\mathcal{R}$ is decreased.

One may wonder whether a more efficient algorithm such as the PPA can yield much better results. In section 4 we take into account that each reset step takes a finite time (without discussing the value of $\mathcal{R}$). We employ entropy considerations (following [4]) that provide lower bounds on the number of reset steps required (ideally) for the PPA and hence also for any cooling algorithm. This result clarifies that there are limitations on AC for *any* cooling algorithm, which makes one wonder whether AC is useful for improving the SNR in magnetic resonance spectroscopy.

In section 5 we present two important cases in which AC might still be found very useful for that purpose in magnetic resonance spectroscopy for noninvasive clinical diagnosis. For instance, AC may supply the same quality of measurements as the currently used methods but with significantly lower amounts of RF radiation applied to the patient.

2 Bonacci algorithms

It is interesting to consider the family of cooling algorithms which asymptotically achieve a Fibonacci-like sequence as the final biases [4, 15]. The first example is the Fibonacci algorithm [4, 15, 28] which is recursively defined as

Fibonacci, m-Fibonacci and δ-Fibonacci: Run $\mathcal{F}(n,n)$

$$\mathcal{F}(n,k) = [\mathcal{F}(n,k-1)\ 3BC(k)]^{m_{n,k}}\ \mathcal{F}(n,k-1), \tag{1}$$

The operations are applied from right to left, the MSB (n) is designated to attain the highest bias, $3BC(k)$ is the 3-bit compression of bits k, $k-1$ and $k-2$, implemented by the gate $|100\rangle \leftrightarrow |011\rangle$ (the MSB is on the left), and $\mathcal{F}(n,2)$ is a RESET step on the two LSBs (bits 1 and 2).

In the Fibonacci algorithm, the number $m_{n,k}$ of iterations can be freely chosen. In the m-Fibonacci algorithm [28], $m_{n,k} = m$ is fixed (for all n,k). For example, with $m_{n,k} = 2$, the algorithm 2-Fib (or 2Fib) acting on n spins, has this subroutine in it: $\mathcal{F}(n,4) = \mathcal{F}(n,3)3BC(4)\mathcal{F}(n,3)3BC(4)\mathcal{F}(n,3)$. After another step in the recursion, we get

$$\begin{aligned}\mathcal{F}(n,4) &= \mathcal{F}(n,2)3BC(3)\mathcal{F}(n,2)3BC(3)\mathcal{F}(n,2)3BC(4)\\ &\cdot\ \mathcal{F}(n,2)3BC(3)\mathcal{F}(n,2)3BC(3)\mathcal{F}(n,2)3BC(4)\\ &\cdot\ \mathcal{F}(n,2)3BC(3)\mathcal{F}(n,2)3BC(3)\mathcal{F}(n,2).\end{aligned}$$

For infinite m the resulting biases of cooled spins are exactly the Fibonacci series $\{\ldots,13,8,5,3,2,1,1\}$.

Regarding $\mathcal{F}(n,2)$ we consider (in all the Fibonacci algorithms) two cases:

- Both LSBs are reset bits: $\mathcal{F}(n,2) = WAIT$
- Only bit 1 is a reset bit: $\mathcal{F}(n,2) = \mathcal{F}(n,1)2BC(2)\mathcal{F}(n,1)$, where $2BC$ ($|10\rangle \leftrightarrow |01\rangle$) is SWAP and $\mathcal{F}(n,1) = WAIT$.

Clearly the second case simply doubles the number of reset steps.

In the algorithm *delta*-Fibonacci, the number of iterations $m_{n,k}$ is chosen so that the bias of bit k will become at least $F_k\left(1-\delta^{n-(k-1)}\right)\varepsilon_0$, where F_k is the k^{th} number in the Fibonacci sequence. For $\delta = \frac{1}{2}$ the required number of iterations $m_{n,k}$ is bound from above [4, 28] by $n-k+2$. The MSB will then attain a bias increase factor of at least $F_n/2$, where the runtime of the algorithm will be bound by $n!$ [4, 28].

For example if we take 4 bits, and $\delta = \frac{1}{2}$, the goal will be $\left\{1.5, 1.5, \frac{7}{8}, \frac{15}{16}\right\}$. The algorithm $\mathcal{F}(4,3)$ will be as follows:

$$\begin{aligned}\{0,0,0,0\} \xrightarrow{\mathcal{F}(4,2)} \{0,0,1,1\} &\xrightarrow{\mathcal{F}(4,2)3BC(3)} \{0,1,1,1\}\\ &\xrightarrow{\mathcal{F}(4,2)3BC(3)} \{0,1.5,1,1\}\\ &\xrightarrow{\mathcal{F}(4,2)3BC(3)} \{0,1.75,1,1\}\\ &\ldots\end{aligned} \tag{2}$$

Notice that the goal has been reached already at the second iteration, $m_{4,3}$=2, if we skip the third iteration, then $\mathcal{F}(4,4)$ with $m_{4,4} = 2$ will be as follows:

$$\begin{aligned}\{0,0,0,0\} \xrightarrow{\mathcal{F}(4,3)} \{0,1.5,1,1\} &\xrightarrow{\mathcal{F}(4,3)3BC(4)} \{1.25,1.5,1,1\}\\ &\xrightarrow{\mathcal{F}(4,3)3BC(4)} \left\{1\frac{7}{8},1.5,1,1\right\}\end{aligned} \tag{3}$$

Which is already beyond the goal. If we choose $m_{4,3} = 3$ as the bound allows, then $\mathcal{F}(4,4)$ with $m_{4,4} = 2$ will be as follows:

$$\begin{aligned}\{0,0,0,0\} \xrightarrow{\mathcal{F}(4,3)} \{0,1.75,1,1\} &\xrightarrow{\mathcal{F}(4,3)3BC(4)} \left\{1\frac{3}{8},1.75,1,1\right\}\\ &\xrightarrow{\mathcal{F}(4,3)3BC(4)} \left\{2\frac{1}{16},1.75,1,1\right\}\end{aligned} \tag{4}$$

For the case of 4 bits it is also possible to run Tribonacci:
Tribonacci: Run $\mathcal{T}(n,n)$

$$\mathcal{T}(n,k) = [\mathcal{T}(n,k-1)\ 4BC(k)]^{m_{n,k}}\ \mathcal{T}(n,k-1), \tag{5}$$

Where $4BC$ is implemented by $|1000\rangle \leftrightarrow |0111\rangle$, $\left(\text{which for biases} \ll 1 \text{ achieves } \varepsilon_4 \leftarrow \frac{3\varepsilon_4+\varepsilon_3+\varepsilon_2+\varepsilon_1}{4}\right)$, and $\mathcal{T}(n,3) = \mathcal{F}(n,3)$ is Fibonacci on the three LSBs, run enough times to reach the desirable bias. In the case of 4 bits, the algorithm is simply $4BC(4)$ and $F(4,3)$ iteratively and the 4-term Tribonacci sequence {4,2,1,1} is asymptotically approached. Similar to Fibonacci, we can choose a goal of $T_k(1-\delta^{n-(k-1)})\varepsilon_0$ for each bit k, where T_k is the k^{th} Tribonacci number. So the goal for 4 spins will be $\left\{2, 1.5, \frac{7}{8}, \frac{15}{16}\right\}$. The values of $m_{n,k}$ satisfying this goal, with $\delta = \frac{1}{2}$ can be chosen to be $m_{4,3} = 2$ and $m_{4,4} = 3$.

$$\begin{aligned}\{0,0,0,0\} \xrightarrow{\mathcal{F}(4,3)} \{0,1.5,1,1\} &\xrightarrow{\mathcal{F}(4,3)4BC(4)} \left\{\frac{7}{8},1.5,1,1\right\} \\ &\xrightarrow{\mathcal{F}(4,3)4BC(4)} \left\{1\frac{17}{32},1.5,1,1\right\} \\ &\xrightarrow{\mathcal{F}(4,3)4BC(4)} \left\{2\frac{3}{128},1.5,1,1\right\}\end{aligned} \tag{6}$$

We can also choose $m_{4,3} = 3$ and $m_{4,4} = 2$ in order to compare it to our calculations of δ-Fibonacci:

$$\begin{aligned}\{0,0,0,0\} \xrightarrow{\mathcal{F}(4,3)} \{0,1.75,1,1\} &\xrightarrow{\mathcal{F}(4,3)4BC(4)} \left\{\frac{15}{16},1.75,1,1\right\} \\ &\xrightarrow{\mathcal{F}(4,3)4BC(4)} \left\{1\frac{41}{64},1.75,1,1\right\} \\ &\xrightarrow{\mathcal{F}(4,3)4BC(4)} \left\{2\frac{43}{256},1.75,1,1\right\}\end{aligned} \tag{7}$$

where we provide also $m_{4,4} = 3$ for comparison.

We see that for the same $m_{n,k}$ values used for δ-Fibonacci, the Tribonacci algorithm not only fails to reach its goal, it even yields lower values than Fibonacci algorithm. Clearly, something is wrong in the definition of the Tribonacci algorithm as is given in [15, 16]. This problem becomes even more severe when looking at high-term bonacci algorithms of [15, 16], such as 4-term-bonacci, etc., and finally k-term-bonacci and All-bonacci. In order to identify the source of the problem it is best to look back into the Fibonacci algorithm.

We can easily identify that the problem exists also in the Fibonacci algorithm: We notice that at the first iteration of $\mathcal{F}(4,4)$ in Eq. 3, replacing the 3BC with $2BC(4)$ will attribute bit 4 with a higher bias

$$\begin{aligned}\{0,0,0,0\} \xrightarrow{\mathcal{F}(4,3)} \{0,1.5,1,1\} &\xrightarrow{\mathcal{F}(4,3)2BC(4)} \{1.5,1.5,1,1\} \\ &\xrightarrow{\mathcal{F}(4,3)3BC(4)} \{2,1.5,1,1\}\end{aligned} \tag{8}$$

Following Eq. 4 and setting $m_{4,3} = 3$, we will achieve an even higher bias

$$\begin{aligned}\{0,0,0,0\} \xrightarrow{\mathcal{F}(4,3)} \{0,1.75,1,1\} &\xrightarrow{\mathcal{F}(4,3)2BC(4)} \{1.75,1.75,1,1\} \\ &\xrightarrow{\mathcal{F}(4,3)3BC(4)} \{2.25,1.75,1,1\}\end{aligned} \tag{9}$$

Notice that in both Eqs. 8 and 9 the goal has been achieved already after a single iteration. When running the Fibonacci algorithm on more than 3 bits replacing $3BC(k)$ by $2BC(k)$ in the first iteration of each recursive level will increase the final bias, and the larger k is — the larger is the gain relative to the original Fibonacci algorithm. Alternatively the same bias goal of $F_k\left(1-\delta^{n-(k-1)}\right)\varepsilon_0$ can be achieved with less iterations. We name this version of the algorithm *new-Fibonacci*.

An interesting example is when we choose a constant number of iterations, $\forall\{n,k\}$, $m_{n,k} = 2$. In this case the original Fibonacci algorithm will become the algorithm 2Fib, as earlier mentioned, and the (better) new Fibonacci, will become new-2Fib, identical to an algorithm named PAC3 [28]. So from what we learn here about old vs. new Fibonacci we conclude that PAC3 always outperforms 2Fib.

Similar to new-Fibonacci, we also define new-Tribonacci: it is wiser to use $2BC(k)$ instead of $4BC(k)$ in the first iteration of each recursive level, and in the next iteration(s) $3BC(k)$ will give the highest bias to bit k. Choosing $m_{4,3} = 2$ we get

$$\begin{aligned}\{0,0,0,0\} \xrightarrow{\mathcal{F}(4,3)} \{0,1.5,1,1\} &\xrightarrow{\mathcal{F}(4,3)2BC(4)} \{1.5,1.5,1,1\} \\ &\xrightarrow{\mathcal{F}(4,3)3BC(4)} \{2,1.5,1,1\} \\ &\xrightarrow{\mathcal{F}(4,3)4BC(4)} \{2.5,1.5,1,1\}\,.\end{aligned} \tag{10}$$

And choosing $m_{4,3} = 3$ we get:

$$\begin{aligned}\{0,0,0,0\} \xrightarrow{\mathcal{F}(4,3)} \{0,1.75,1,1\} &\xrightarrow{\mathcal{F}(4,3)2BC(4)} \{1.75,1.75,1,1\} \\ &\xrightarrow{\mathcal{F}(4,3)3BC(4)or4BC(4)} \{2.25,1.75,1,1\} \\ &\xrightarrow{\mathcal{F}(4,3)4BC(4)} \left\{2\frac{5}{8},1.5,1,1\right\}.\end{aligned} \tag{11}$$

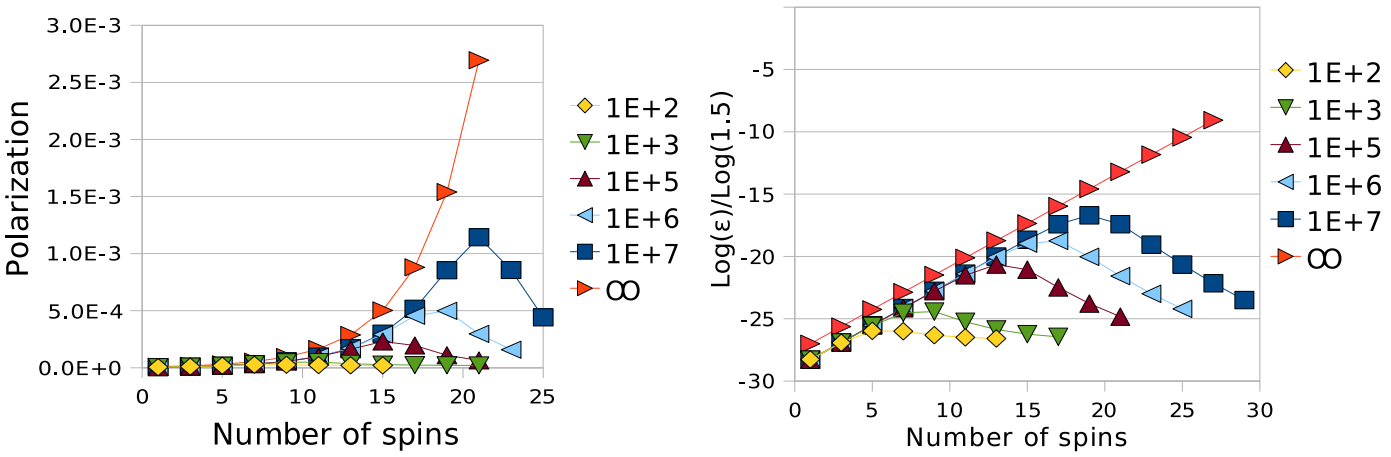

Figure 1: Simulation results: the diagram on the left shows the maximal bias achievable by 2PAC on spin systems with various $\mathcal{R}$ values ranging from 100, the scale which may be attainable in liquid-state NMR, to 10^6 attainable in solid-state NMR using spin-diffusion as the reset mechanism; where the red graph is the ideal case of $\mathcal{R} = \infty$. The equilibrium bias is set to $\varepsilon_0 = 10^{-5}$. We see that even for the enourmous $\mathcal{R} = 10^7$, bias increase by 2PAC is not possible for more than two orders of magnitude and for no more than 20 spins. The diagram on the right is the same diagram where the y axis is in logarithmic scale, which shows us that the cooling is exponential until reaching the highest achievable bias, even for the smallest $\mathcal{R}$ simulated.

Notice that for both Eqs. 10 and 11 the goal has already been reached by the 2^{nd} iteration, before we even used a 4BC gate. This potentially reduces the time complexity for running Tribonacci or All-bonacci [15, 28] on larger spin systems.

We define new-Tribonacci to be better then new-Fibonacci, by chosing the optimal compression gate (2BC, 3BC, or 4BC) at each step. Similarly, new-k-bonacci is therefore better than new-{k-1}-bonacci, by definition.

It is interesting to research whether the new-δ-bonacci algorithms can converge to the goal as defined earlier within a reasonable time. We know that the time complexity of the old version of δ-fibonacci is $O(n!)$, hence, by definition, any new-bonacci will surely reach the Fibonacci goal, but what is the complexity of new-δ-Tribonacci for reaching the Tribonacci goal? And similarly, for the other k-bonacci algorithms? And how significant is the advantage of the new versions of the algorithms? These questions are left open for future research.

Another open research question arises from the redundancy in the algorithms. E.g., after the first $2BC(4)$ step in new-Tribonacci, the next subroutine is $\mathcal{F}(4,3)$ where the first step is a redundant reset $\mathcal{F}(4,2)$ because the $2BC(4)$ step did not destroy the two reset spins anyhow. So clearly it is interesting in practice (and maybe also possible in theory) to optimize the algorithm further by removing such redundant steps.

3 Accounting for relaxation

The definition of mPAC is

$$\begin{aligned} M_j(k) &= [\mathcal{B}(k)M_{j-1}(k-2)\ PT(k-2 \to k-1) \\ &\quad M_{j-1}(k-2)]^m\ PT(k-2 \to k)\ M_{j-1}(k-2). \end{aligned} \tag{12}$$

The polarization transfer (PT) moves the bias of the first spin onto the second spin, such that $PT(A \to B)$ assigns to spin B the bias of spin A. For the purpose of this work, whenever PT is used, it is identical to SWAP as we do not care about the polarization of spin A after the PT or SWAP operation; and the advantage of PT is that it is easier to implement in practice.

For example, using the genral formula above, 2PAC on 3 spins is defined as

$$M_1(3) = [3BC(3)M_0(1)\ PT(1 \to 2)\ M_0(1)]^2\ PT(1 \to 3)\ M_0(1), \tag{13}$$

where $M_0(1)$ is a reset step of the reset bit (spin 1). For 5 spins 2PAC is defined as follows:

$$M_2(5) = [3BC(5)M_1(3)\ PT(3 \to 4)\ M_1(3)]^2\ PT(3 \to 5)\ M_1(3). \tag{14}$$

The algorithm under this definition cools only the MSB. In order to cool all the bits from 1 to k we add the cooling of all less significant bits in decreasing order. For example

$$M_{all}(3) = M_{all}(1)\ PT(1 \to 2)\ M_0(1)\ M_1(3) \tag{15}$$

$$M_{all}(5) = M_{all}(3)\ PT(3 \to 4)\ M_1(3)\ M_2(5) \tag{16}$$

$$M_{all}(k) = M_{all}(k-2)PT(k-2 \to k-1)\ M_{j-1}(k-2)\ M_j(k), \tag{17}$$

where $M_{all}(1) = M_0(1) = WAIT$. The simulation results we show here are of $M_{all}(7)$.

We shall now deal with finite $\mathcal{R}$, and analyze 2PAC, assuming one reset spin, for several values of $\mathcal{R}$. Let $\mathrm{T}_{\mathrm{WAIT}} = d \cdot \mathrm{T}_1(\mathrm{reset})$ be the duration of each reset step, namely, we decide here that the time we WAIT and do nothing between computing steps of the algorithm, is always the same. When $d \approx 5$, the polarization bias of the reset spin approaches

its equilibrium value ε_0 by the end of each reset step. We also want $\mathrm{T_{WAIT}} \ll \mathrm{T_1(comp)}$ in order to enable many reset steps in our algorithm. Ignoring the duration of unitary transformations, if an algorithm has N_r reset steps, each requiring a waiting time $\mathrm{T_{WAIT}}$, then the runtime of the algorithm is $\mathrm{T_{run}} = \mathrm{T_{WAIT}} \cdot N_r = d \cdot N_r \cdot \mathrm{T_1(reset)}$. For given $\mathrm{T_{WAIT}}$ and $\mathrm{T_1(comp)}$, it is desired that the runtime satisfies $\mathrm{T_{run}} = \mathrm{T_{WAIT}} \cdot N_r \ll \mathrm{T_1(comp)}$, hence we get a condition on the number of reset steps N_r. When the runtime satisfies $\mathrm{T_{run}} \ll \mathrm{T_1(comp)}$, and $d = 5$, the ideal results obtained in [28] still apply with only minor corrections. However, when $\mathrm{T_1(comp)}/\mathrm{T_{run}}$ is in the range of about 0.1 to 10, namely, $\mathcal{R}/(d \cdot N_r)$ is between 0.1 and 10, the deviation from the ideal result is significant, while for $\mathcal{R}/(d \cdot N_r) > 10$ the deviation is still small. Although this is redundant, it is still useful to define another dimensionless parameter $\mathcal{Q} = \mathrm{T_1(comp)}/\mathrm{T_{WAIT}}$, such that $1/\mathcal{Q}$ provides the extent by which the computation spins relax during one WAIT period.

In order to run a long algorithm (with many reset steps) and still reach near-ideal cooling, we would want the durations to satisfy $\mathrm{T_1(comp)} \gg \mathrm{T_{run}} \gg \mathrm{T_{WAIT}} \gg \mathrm{T_1(reset)}$. It is thus natural to define another (redundant) dimensionless parameter, $\mathcal{D} = \mathrm{T_1(comp)}/\mathrm{T_{run}}$ (hence $\mathcal{D} = \mathcal{Q}/N_r$), and now $\mathrm{T_1(comp)} = \mathcal{D}\mathrm{T_{run}} = \mathcal{D}N_r\mathrm{T_{WAIT}} = \mathcal{D}N_r d\mathrm{T_1(reset)}$, hence $\mathcal{R} = \mathcal{D}N_r d$. Note that as long as $\mathcal{D} \gg 1$ (namely, $\mathcal{Q} \gg N_r$) the computation spins barely relax, and as long as $d \gg 1$ the reset spin regains a bias very close to its equilibrium bias.

Using a constant $\mathrm{T_{WAIT}}$ in our calculations allows a major simplification: We can simply replace the equilibrium reset-spin bias by $(1-e^{-d})\varepsilon_0$, assuming the reset spin has no polarization prior to each reset step. For each computation spin, one might expect to replace the ideal final bias, ε_f by $e^{-(N_r/\mathcal{Q})}\varepsilon_f$ (namely, $e^{-(1/\mathcal{D})}\varepsilon_f$), which would be true if that spin got its final optimal bias at the beginning of the relaxation process. However, the bias is gradually increased by portions that are subject to less relaxation, hence the above result provides a lower bound on the final polarization.

For n spins the calculations are actually rather cumbersome, as there are 2^n terms in the diagonal density matrix, at each step. Taking into account finite d and $\mathcal{Q}$ becomes non-trivial, hence we take here an *extended-Markovian* approach: after each reset step, not only does the environment "forget" any remaining correlations with our system, but also the reset spin is considered as a separate "environment" and becomes uncorrelated with the computation spins. For $\mathcal{R} \to \infty$, calculations using this model become identical to previous calculations done for ideal theoretical AC [1–4, 15, 16, 18, 28, 32], if both d and $\mathcal{Q}$ go to infinity. When d is finite, and we use this model, the bias of the reset spin evolves from some initial value ε into $(\varepsilon - \varepsilon_0)e^{-d} + \varepsilon_0$. Already for $d = 5$ the extended-Markovian approach seems reasonable, as can be verified (numerically) for small values of n. However, it might be much less accurate for larger values of n, if compared with a model taking all correlations into account. The extended Markovian model influences the cooling process by directly removing only the correlations of the reset spin with the computer spins. However, for mPAC, this then indirectly removes undesired correlations among the computer spins as well.

Using this model, we simulated 2PAC with 7 spins, starting from $\varepsilon_0 \ll 1$, with reset steps satisfying $d = 5$. Cooling all the spins, see [28] section IV, requires $N_r = 187$ steps, and ideally attributes the most significant bit (MSB) with a final bias of 5.36 (in units of ε_0), see Table II in [28]. This cooling level remains, as long as $\mathcal{D} \gg 1$; When $\mathcal{R} = 10000$ (i.e. $\mathcal{Q} = 2000$), the MSB attains a final bias of 5.11. When $\mathcal{D} \sim 1$, the cooling level is significantly reduced; $\mathcal{R} = 1000$ (namely $\mathcal{Q} = 200$) leads to a final bias of 3.63. Finally, when $\mathcal{D} \ll 1$, the cooling becomes quite negligible; When $\mathcal{R} = 100$ (namely $\mathcal{Q} = 20$), the final bias is only 1.07. Better results could be achieved if the two LSBs are reset bits. Another way to reach better results is to cool only the MSB. Notably, $\mathcal{R} \sim 10000$ was achieved in solid-state NMR, using spin-diffusion for rapid repolarization [29, 30], but even $\mathcal{R} = 100$ seems beyond the capability of current liquid-state NMR.

The formal proof that the consequence of the extended Markovian model for the mPAC algorithm is the removal of all the spin correlations is not given here with full details, but let us consider an illustrative example of using 2PAC on the four LSB of a spin system with four or more spins, where only A is a reset spin. Suppose that at some stage in the process of cooling, the 3 spins have a known bias configuration $\{\varepsilon_C, \varepsilon_B, \varepsilon_A\}$, such that the spins are potentially correlated among them and/or with other spins. In the extended Markovian model, cooling of these 3 spins, by repeated reset steps and SWAPs, removes their correlations, as we now show. We start after spin A had been reset to some near-equilibrium — after a WAIT of a duration of $5 \cdot T_1(A)$, the bias of A will be $\varepsilon_r = (\varepsilon_A - \varepsilon_0)e^{-5} + \varepsilon_0$, with no remaining correlations with the other spins, due to the extended-Markovian assumption. The reset of the three spins now goes as follows:

$$\left\{\underline{\varepsilon_C}, \underline{\varepsilon_B}, \varepsilon_r\right\} \xrightarrow{A \leftrightarrow C} \left\{\varepsilon_r, \underline{\varepsilon_B}, \underline{\varepsilon_C}\right\} \xrightarrow{WAIT} \left\{\varepsilon_{r'}, \underline{\varepsilon_B}, \varepsilon_r\right\} \xrightarrow{A \leftrightarrow B} \left\{\varepsilon_{r'}, \varepsilon_r, \underline{\varepsilon_B}\right\} \xrightarrow{WAIT} \left\{\varepsilon_{r''}, \varepsilon_{r'}, \varepsilon_r\right\}, \tag{18}$$

where $\varepsilon_{r'} = (\varepsilon_r - \varepsilon_0)\, e^{-\frac{T_{WAIT}}{T_1(comp)}} + \varepsilon_0$, and $\varepsilon_{r''}$ is similar but with $2T_{WAIT}$, and where we underline each spin when it is still potentially correlated. We continue according to mPAC (see Eq. 12) for $m = 2$, applied onto 4 spins, $DCBA$, where only A is a reset spin and the bias configuration is $\left\{\varepsilon_D, \underline{\varepsilon_C}, \underline{\varepsilon_B}, \underline{\varepsilon_A}\right\}$, as is typical after a 3BC on bits CBA. Here ε_D is unknown and D is not correlated to any other spin(s). We underline the potentially correlated spins. The next step of 2PAC is using SWAP steps between D and the cooled C leading to the configuration$\left\{\underline{\varepsilon_C}, \varepsilon_D, \underline{\varepsilon_B}, \underline{\varepsilon_A}\right\}$ After a WAIT of a duration of $5 \cdot T_1(A)$, the bias of A will be $\varepsilon_r = (\varepsilon_A - \varepsilon_0)e^{-5} + \varepsilon_0$ with no remaining correlations with the other spins[2]. The extended Markov assumption means that spin A will now be uncorrelated to the rest of the spins on the molecule, and therefore in a tensor product state with all other spins. The following steps of 2PAC are:

$$\left\{\underline{\varepsilon_C}, \varepsilon_D, \underline{\varepsilon_B}, \varepsilon_r\right\} \xrightarrow{A \leftrightarrow C} \left\{\underline{\varepsilon_C}, \varepsilon_r, \underline{\varepsilon_B}, \varepsilon_D\right\} \xrightarrow{WAIT} \left\{\underline{\varepsilon_C}, \varepsilon_r, \underline{\varepsilon_B}, \varepsilon_r\right\} \xrightarrow{A \leftrightarrow B} \left\{\underline{\varepsilon_C}, \varepsilon_r, \varepsilon_r, \underline{\varepsilon_B}\right\} \xrightarrow{WAIT} \left\{\varepsilon_C, \varepsilon_r, \varepsilon_r, \varepsilon_r\right\}. \tag{19}$$

It is easily possible to prove now, by induction, that for any number of spins, correlations are completely removed prior to all compression steps.

[2] In the simulation we chose $\varepsilon_A = 0$, and so $\varepsilon_r = (1 - e^{-5})\varepsilon_0$. The computation spins also relax during the WAIT, but we ignore this here, dealing only with the issue of removal of correlations. In the simulation we took all relaxations into consideration.

4 Lower bounds on required reset steps

The previous section dealt with a specific practicable algorithm. Could much better results be obtained using a much more efficient algorithm, or even the optimal algorithm (PPA)?

AC relies on thermalization (or other means of repolarization) for reset operations following selective PT from a reset spin to the target spin or following polarization compression. Ideally, optimal AC asymptotically cools one computation spin in an n-spin system (with one reset spin) by a factor of 2^{n-2}, when the final polarization is $\lesssim 1\%$. The optimal PPA algorithm was proven to approach this limit, and to require a minimal number of reset steps at each stage [3, 4]. As in the previous section, also here we neglect the time of the unitary evolution in between reset steps. The use of an extended-Markov model cannot be justified for the PPA as correlations among computation spins are not removed by reset steps followed by SWAPs, etc. Thus, we analyze here the ideal case with infinite $\mathcal{Q}$. If d is infinite then this is exactly the ideal model, and if d is finite then only the bias ε should be modified accordingly. Since we focus on infinite $\mathcal{Q}$, the limitations here are not as novel as the ones discussed in the above section, but are simple consequences of our earlier results together with Schulman [3, 4]. Such results, when written the way they appear here, simplify the comparison of AC with other methods used to improve the SNR in magnetic resonance spectroscopy.

We now present lower bounds (based on entropy considerations) on the number of reset steps required by any spin-cooling algorithm [3, 4].

Theorem 1 *Given an n-spin system with a single reset spin, and $n-1$ computation spins, where all spins are in a completely mixed state and the polarization bias of the reset spin after a reset step is $\varepsilon \ll 1$, the lower bound on the number of reset steps required to cool one computation spin to $k\varepsilon \ll 1$ is k^2.*

Proof:

Our goal is to cool one spin to $k\varepsilon$, whereby the total entropy of the spin system, $H_{init} = n$, would decrease to $H_{fin} \leq n-(k\varepsilon)^2/\ln 4 = n-k^2\varepsilon^2/\ln 4$; the equality is obtained when the other $n-1$ spins remain completely mixed (no other spin is cooled). The entropy removed from the entire spin system is therefore $\Delta H = H_{init} - H_{fin} \geq \left(k^2\right)\varepsilon^2/\ln 4$. Each reset step reduces the entropy by at most [8] the information content of a single reset-spin at thermal equilibrium $I_1 = \varepsilon^2/\ln 4$ (when the reset spin goes from the completely mixed state to equilibrium). Therefore the lower bound on the number of reset steps is given by: $\frac{\Delta H}{I_1} \geq \frac{\left(k^2\right)\varepsilon^2/\ln 4}{\varepsilon^2/\ln 4} = k^2$. □

¿From this theorem it follows that for optimal cooling [3], where $k = 2^{n-2}$, the lower bound on the number of required reset steps is about k^2 already for small spin systems, since $k^2 = 2^{2(n-2)} \gg n$ is satisfied already for $n \gtrsim 6$. It is therefore much more efficient (in terms of average entropy reduction per reset step) to stop the PPA after much fewer reset steps.

5 Prospects of algorithmic cooling and heat-bath cooling for NMR spectroscopy

We investigated the feasibility of applying the heat-bath cooling (HBC) experiments, as well as algorithmic cooling (AC), to improve NMR spectroscopy. We emphasize in vivo ^{13}C spectroscopy (13C-MRS) in the brain, where ^{13}C-labeled metabolites are observed over extended periods of time (up to several hours). One of the main targets of in vivo brain spectroscopy has been glutamate (Glu), a major excitatory neurotransmitter, to which we applied HBC [10]. Despite focusing on 13C-MRS, the results here, especially those in section 5.3, are more general.

5.1 Carbon-based brain spectroscopy

The application of NMR spectroscopy in vivo, for non-invasive study of metabolic processes in living organisms, is called *magnetic resonance spectroscopy (MRS)*. Proton (^{1}H) MRS is applied clinically since the late 1980s, and was approved by the US Food and Drug Administration (FDA) about a decade ago; It is commonly performed on conventional MR Imaging (MRI) systems. In the brain, proton MRS measures about 35 metabolites, a small portion of the estimated 2,000-20,000 metabolites [33]. Localization to a specific volume in the brain allows to focus on the metabolism in a small region of interest, such as a malignant tissue. Proton MRS is also applied clinically for diagnosis of metabolic disorders, where abnormal enzymatic activity produces unusual metabolite concentrations [34]. MRS of the living brain provides a "virtual biopsy" [33], revealing information not accessible by other means.

One main goal of (brain) MRS is to provide useful diagnostic information for clinical investigation; for this purpose, the steady-state concentration measurements obtained by *static analysis* with proton MRS are quite limited [35]. While proton MRS is very common, carbon is the ideal element for following metabolic processes; it is present in most metabolites and may be labeled by enriching the stable isotope ^{13}C from its natural abundance of 1.1% up to $\sim 99\%$. The label can be followed dynamically, over extended periods of time, with negligible background, since the major isotope (^{12}C) lacks magnetic moment. In contrast, hydrogen-based MRS is not useful for following dynamic processes, since both stable isotopes are NMR-active, and the much less abundant ^{2}H isotope (spin-1) has very low sensitivity and small chemical shifts (about six times smaller than for ^{1}H).

13C-MRS is developed since the 1970s and was shown in the early 1990s to detect labeled glutamate (Glu) in the human brain [36] following infusion of ^{13}C-labeled glucose. It provides enhanced resolution over proton MRS, due to the much wider (by over 20-fold) chemical shift range of ^{13}C [37, 38]. However, its SNR is significantly lower than proton MRS, such that many scans are generally needed for a clinically-useful diagnosis. While technically challenging, 13C-MRS has been established as the only noninvasive method for measuring neurotransmission by Glu and the energetics of specific cells in the human brain [39]. It has provided rates for specific metabolic routes, such as neurotransmitter cycling (Glu/glutamine) [35], allowed to establish the energy cost of brain function, and provided

various other highly important information (reviewed in Ref [39]). Recently, 13C-MRS also provided exchange rates for individual enzymes [40], using saturation transfer.

13C-MRS supports *personalized medicine*[3] and may provide initial diagnosis, monitor treatment efficacy, and determine optimal treatment and dosage [33]. In contrast, proton MRS is not useful for in vivo drug monitoring, since contrast agents either do not enter the brain, or they interact with tumors differently than the drug. While still pending FDA approval for routine clinical usage, 13C-MRS is showing promising clinical potential; recently, by acquiring spectra briefly at steady state after infusion of labeled acetate, 13C-MRS enabled rapid assessment of inflammation in glia cells, indicative of neurodegeneration [41]. More recently, this protocol was applied to patients of Alzheimer's disease and showed an increase in glial metabolism, which was correlated with cognitive decline [42]. Several major challenges remain in the path of clinical 13C-MRS, including (mainly) low SNR, which requires acquisition of relatively large volumes (several mL), and also tissue heating and prolonged infusion periods (typically 2 hours) [39].

By employing hyperpolarized substrates, the SNR of 13C-MRS can be dramatically improved, enabling metabolic imaging [43]. Human clinical trials were recently conducted and showed promising results for prostate cancer using dissolution dynamic nuclear polarization with hyperpolarized ^{13}C-pyruvate [44]. Since the hyperpolarization decays after few minutes at most, applicability to brain spectroscopy is limited [45].

5.2 AC for ^{13}C brain spectroscopy

AC requires metabolites where there is a string of qubits, with at least one reset spin. The portion of such metabolites may be increased by using a suitable *isotopomer* (isomer with specific labeling). For example, fully-labeled glucose might be a better substrate than the single or double labeled isotopomers. Recently, metabolism of different substrates, e.g. glucose and lactate, was detected simultaneously at high field (11.7T), based on different line-splitting patterns [46]. Heteronuclear PT from ^{1}H to ^{13}C, which may be considered as a relevant starting point for AC and HBC, was successfully implemented in the brain in a recent in vivo study of aging [47].

Theoretically, for n spins, AC was shown to produce exponential cooling of $\left(\frac{3}{2}\right)^{\lfloor (n-1)/2 \rfloor}$ via the most basic algorithm [2], and up to 2^{n-2} [3, 4, 15] via the optimal algorithms. However, in practice, cooling is limited by various physical constraints, which include experimental relaxation/repolarization rates (as shown in section 3), decoherence, imperfect pulses, and RF inhomogeneity. Moreover, even in the ideal case, we show in section 5.3 that, contrary to initial expectations [2, 15], for improving the SNR of a particular spin using one reset spin, it is often more efficient to perform multiple repetitions of PT than to use AC. It may, however, still be advantageous to perform AC in several cases; There are three important areas where AC could contribute to NMR spectroscopy: when *multiple reset spins* are available (which we do not discuss here); *spectral editing* to resolve overlapping signals (section 5.5); and reduction of the *specific absorption rate (SAR)*, which is highly important for in vivo carbon-based brain spectroscopy (section 5.4).

5.3 Sensitivity enhancement - AC vs signal-averaging

It is interesting to note that the use of k^2 reset steps as in theorem 1, is also highly relevant in a more conventional approach in NMR: signal averaging. By averaging over $N_r = k^2$ scans, after PT from a reset spin the SNR is also increased by a factor of k, hence the advantage of AC seems to be lost. However, the significantly fewer scans required by AC might be advantageous in reducing the specific absorption rate (SAR).

AC relies on thermalization (or other means of repolarization) for reset operations. It is therefore natural to compare its performance with multiscan-PT, a form of signal averaging that enhances the SNR by performing several scans, each following selective PT from a highly-polarized reset spin to the target spin, with a delay between scans to allow the reset spin to repolarize. Here we ignore the high polarization of the reset spin, which for protons and carbons simply adds a multiplicative factor of four to both AC and signal averaging, and we focus on the short reset time of the reset spin. In contrast to the conventional multiscan used in signal averaging, the delay between consecutive scans in multiscan-PT is much shorter, as it is determined by the relaxation time of a fast-relaxing spin. Ideally, optimal AC asymptotically cools one computation spin in an n-spin system, where one spin is a reset spin, by a factor of 2^{n-2}, up to final polarizations of about 1% [3]. The optimal *partner pairing algorithm (PPA)* approaches this limit using a minimal number of reset steps. In practice (e.g., if analyzed using the tools introduced in section 3) it is less clear if AC will be sufficiently efficient relative to multiscan-PT.

5.4 AC for reduction of SAR

The specific absorption rate (SAR) measures the rate at which electromagnetic (RF) radiation is absorbed by the body, reflecting the RF power applied during pulse transmission. SAR levels (in W/kg) are tissue dependent, e.g. in the frontal lobe (around the eyes) the SAR is high due to limited circulation [39]. Local SAR levels are obtained by numerical methods, as it is difficult to measure the relevant electrical fields or temperatures in vivo [38]. Alternatively, a three-dimensional sample (phantom) is used, which mimics the target tissue [48]. For in vivo MRI and MRS, SAR levels are tightly regulated; for the entire head, the maximum SAR is 3W/kg averaged over 10 minutes, while for any other tissue, 8W/kg are allowed, averaged over 5 minutes [38].

For 13C-MRS, *proton decoupling*, where the protons are saturated with RF irradiation, is commonly employed during data acquisition to considerably simplify the spectrum (reduce the line splittings of each ^{13}C signal), thereby increasing both SNR and resolution. For example, at 3T the signals of the carboxylic carbon C1 of aspartic acid and Glu largely overlap without decoupling, while they are well-resolved when decoupling is applied (see Fig. 4 in ref [48]). Proton decoupling is typically applied over a prolonged period of time (100 ms to few seconds), and therefore strongly contributes to the SAR [49]. The large heteronuclear scalar coupling ($\sim$ 125–150 Hz) between alkyl carbons (e.g., Glu C4) and their attached protons presents a major problem for in vivo 13C-MRS [38], since effective decoupling requires

[3]People now consider the four "P"s: predictive, personalized, preventive, participatory [33].

very high RF field strength ($\gamma B_2 \gg J_{CH}$, where B_2 is the decoupler field) [50]; In order to meet SAR limits, decoupling is generally limited to surface or half-volume coils rather than the state-of-the-art volume coils. The problem is much worse at higher fields; the decoupling bandwidth grows linearly with magnetic field strength, and the SAR grows quadratically. Already at 3T, SAR limitations considerably reduce the effectiveness of WALTZ-4, the most common in vivo decoupling method that performs well at conditions of low RF homogeneity [48, 50]. Indeed, concerns were raised that the inability to decouple effectively might offset the benefits of higher magnetic fields commercially available.

In 2007, a strategy was devised [50] to overcome the limitation on decoupling described above. By using [2-^{13}C]glucose, rather than the commonly used 1-labeled isotopomer, the label was found to be incorporated in the carboxylic carbons (mainly Glu C5, but also Glu C1 and glutamine C1 and C5 [51]). Even though no protons are directly attached, decoupling is still needed in most cases, although much lower power is sufficient. The SNR of carboxylic/amide carbons is reduced due to their long T_1 and the inability to perform PT (since protons are unavailable). However, the absence of interfering lipid signals opens the possibility of parallel imaging using multiple ^{13}C receivers and whole brain low power decoupling [52]. Further reduction in decoupling power was achieved by using stochastic decoupling, originated by Ernst, which is based on random noise [50]. Such decoupling provided in 2008 acceptable SAR levels at 1.5T in the frontal lobe in human subjects, which is associated with memory and other higher functions [51].

We have seen that when only one reset spin is available, multiscan-PT offers better SNR improvement than AC. However, several acquisition periods (one per scan) are required by multiscan PT, while AC only employs a single acquisition period at the end of the cooling algorithm. Consequently, AC could considerably reduce the SAR. For example, [1,2-$^{13}C_2$]Glu produced in the brain after administration of [1,2-$^{13}C_2$]glucose (see Fig. 5 in Ref [37]) can be cooled (ideally by a factor of 6) by PAC1 (HBC followed by 3B-Comp) with two reset steps. Multiscan PT can cool better (by a factor of $4\sqrt{3} \sim 6.8$), however the two additional proton-decoupled acquisition periods would generate up to three times higher SAR. AC could therefore accommodate the higher power required for standard WALTZ decoupling; Alternatively, the reduced SAR could enable simultaneous acquisition of multiple volume elements (voxels) in the brain, which could allow to compare normal and abnormal brain regions [48] by permitting the use of RF-intensive volume coils.

It is not yet clear whether the reduction in dissipated heat associated with AC would yield a comparative reduction in the SAR. The recent 3T study of the frontal lobe [48] acquired spectra every 4s (irradiating the proton for 205 ms in each scan); during this small inter-scan delay, the slow heat dissipation in the frontal lobe might still limit SAR accumulation over successive scans. The actual effect of AC on the SAR may be estimated using well-established numerical simulations that consider tissue properties (see Ref [48] and references therein).

5.5 AC for spectral editing

We consider *spectral editing* in the broad sense as an operation (pulse sequence) that enhances a selected part of a complex NMR spectrum. In particular, spectral editing techniques (1D and 2D) are useful for selecting subspectra of isotopically-enriched metabolites, such as ^{13}C and ^{15}N-labeled isotopomers [53]. A need for spectral editing is commonly encountered in vivo, where many metabolites are produced. In many cases, several *isotopomers* (isomers labeled at different positions) of the same metabolite are formed, e.g. single-labeled, double-labeled, and triple-labeled Glu obtained after administration of [1,6-$^{13}C_2$]glucose [38, 49].

We [10] applied simultaneous HBC of the two backbone carbons of amino acids, C1 and C2. Both labeled carbons of Glu and Gly were cooled simultaneously by truncated PT via environment thermalization (POTENT) pulse sequences [8], which consisted of a selective-reset (an HCC relay - from the proton to C2 to C1, followed by a reset step) and a second PT from the α protons to C2. This pulse sequence may be used for in vivo spectral editing of glycine isotopomers, to selectively enhance the double-labeled isotopomer; both carbons would be significantly enhanced ($\sim$ 2.5-fold), while for the single labeled [2-^{13}C]Gly isotopomer, only C2 would be enhanced (for [1-^{13}C]Gly, C1 would be diminished due to its long T_1)[4]. For glutamate (and glutamine), similar spectral editing could selectively enhance the 1,2-$^{13}C_2$ *cumomer*[5], which includes all nine isotopomers where both C1 and C2 are labeled, from the other three cumomers produced after infusion of labeled glucose [37, 49].

The pulses applied to the protons are expected to produce limited PT to C3 and C4, since the refocusing delay of the heteronuclear INEPTs is set[6] according to CH [55]. A slight modification of POTENT could prevent any residual polarization enhancement on proton-bearing carbons other than C2, e.g. C3 and C4 of Glu and glutamine. In the improved variant of POTENT, the hard pulses applied to the protons are replaced by spin-selective pulses that target H2. Such pulses may be readily obtained for glutamine and Glu at high field ($\sim$ 10T), where the chemical shift between H2 and the nearby proton (H4) of about 1.5 ppm [53] is sufficient. At lower fields typically used in vivo, selectivity may be obtained by setting the ^{1}H transmitter frequency downfield from H2 (by about 1 ppm). This approach allowed selective addressing of the C4 methylene protons of γ-aminobutyric acid (GABA) [56] at 3T and 7T; in that case, the chemical shift from the adjacent C3 proton was only $\sim$ 1 ppm. The proton-selective POTENT variant is essential for amino acids, which have an additional methine (CH), such as isoleucine, leucine and valine, which have similar chemical shift between H2α and the other protons ($\sim$ 1.5–2 ppm) [53]. Proton signals of all three aliphatic amino acids were observed in a cerebral abscess by ex vivo proton MRS [33].

The original POTENT sequence may be used to cool all four labeled carbons of [1,2,3,4-$^{13}C_4$]Glu, obtained in vivo by metabolism of fully labeled glucose [57]. The four labeled carbons would be cooled to a similar extent by setting the heteronuclear refocusing delay to about $\frac{1}{3J}$ (obtained numerically based on data from ref [53]), such that the efficiency

[4] An HCC relay which enhances only C1 is sufficient to reveal the presence of the double labeled Gly isotopomer.

[5] The concept of a cumulative isotopomer, or cumomer, was introduced about a decade ago to analytically solve nonlinear isotopomer balance equations [54].

[6] Setting the delay to be $\frac{1}{2J}$ would entirely suppress PT from the methylene protons if all J couplings were the same. In practice [49, 53], the stronger scalar coupling of C2 ($\sim$ 145 Hz vs $\sim$ 130 Hz) is expected to allow limited PT to C3 and C4 (at about 30–40% efficiency).

of PT to C2, C3, and C4 becomes similar (ideally about 90%). In this case, the reset step after the HCC relay would apply to all protons due to the similar T_1 relaxation times.

AC may provide spectral editing that resolves the sub-spectrum of a particular spin from overlapping sub-spectra that originate from other spins. While PT from a proton spin can enhance the resolution of the target ^{13}C spin by at most 4-fold, AC could ideally, in the future, attain an additional ten-fold enhancement using 6–7 spins. Such improvement could significantly aid in identifying and quantitating key metabolites that constitute a small percentage of the sample, which would remain masked, after PT, by overlapping signals of other metabolites. A powerful strategy for in vivo spectral editing [56] employed selective addressing and the unique spin-spin coupling of the target spin to remove overlapping signals of several metabolites. For γ-aminobutyric acid (GABA), a chemical shift selective saturation (CHESS) pulse sequence commonly applied (since the mid-1980s) for water suppression[7] [33] suppressed the C4 methylene and its overlapping ^{1}H signals from creatine, glutathione and macromolecules (all around 3 ppm); the target GABA H4 signal was then selectively regenerated by a novel homonuclear PT from the neighboring C3 proton (at 1.91 ppm) which was unaffected by CHESS ($J_{HH} \sim 7$ Hz). AC may provide similar spectral editing for in vivo 13C-MRS by cooling a ^{13}C spin bound to the target spin and transferring the enhanced polarization via PT after suppressing the target carbon and all overlapping signals.

The apparent advantage of AC for spectral editing could be diminished when considering the limited SNR enhancement achieved by AC with respect to multiscan-PT (see section 5.3). Consider the scheme whereby the signal of the target carbon is suppressed and then regenerated from the nearby carbon enhanced by AC. Assuming that a single (proton) reset spin was used, AC is limited with respect to multiscan-PT (from a proton) over the same duration. Yet, AC might still be preferred in practice, due to the lower SAR (see section 5.4). When this scheme is not possible, due to similar coupling patterns for the target spin and overlapping spins (of other metabolites), AC may selectively enhance the target signal beyond PT. However, in this case multiscan-PT grants additional selectivity over the other (carbon) computation spins according to the T_1 ratio, due to the rapid relaxation between scans. The overall selectivity of multiscan-PT is therefore similar to AC, which is limited by the relaxation time ratio [28].

5.6 Discussion of the prospects of AC for MRS

Ideally, AC was shown to reach an SNR improvement that is exponential in the number of spins. However, such cooling is difficult to obtain in practice due to several physical limitations, most notably (for liquid-state NMR) the limited ratios between the relaxation rate of the cooled spins and the repolarization rate of the reset spins [8]. For the case of a single reset spin, a common signal-averaging process, called here multiscan-PT, can improve the SNR of the target spin more than AC over the same period of time. For several reset spins with higher equilibrium polarization, such as protons, one cycle of AC can cool a computation spin, such as a ^{13}C beyond the reset spin polarization. However, the same cooling can in principle be obtained by a more general PT, as shown by Sørensen for I_nS [13] and I_nS_m [58] spin systems. In practice, as noted by Sørensen, experimental implementation of such PT may not be feasible, in particular when the protons are bound to different carbons; in that case, the two steps of PAC (PT and 3B-Comp) may be the only intuitive option[9]. Beyond SNR improvement, AC could contribute to in vivo 13C-MRS in two important areas: spectral editing to resolve complex spectra by selecting for signals of a particular metabolite, and reduction of SAR associated with proton decoupling (by reducing the amount of data acquisitions).

Further research is needed to evaluate both potential applications. For spectral editing, it is necessary to seek practical scenarios where AC and HBC are advantageous. HBC can provide direct spectral indication for the presence of specific (e.g., 1,2-double-labeled) metabolites, while AC generally enhances a single spin considerably, ideally much beyond PT. AC is limited with respect to multiscan-PT, yet it might still be advantageous due to the expected reduction in SAR. For example, AC could enhance spectral editing where the target signal (along with all overlapping signals) is suppressed; the target signal could be regenerated by PT from an adjacent carbon cooled by AC, rather than by PT from a proton. For SAR reduction, the expected effect in vivo can be estimated for multiple scans using well-established numerical simulations.

We hope that our results would encourage further research into experimental open system cooling in liquid NMR. It appears that there is real benefit for *in vivo* heteronuclear brain spectroscopy, a noninvasive method that provides unique information about metabolism in specific compartments of the brain. To this end, further evaluation is needed for the prospective near-future applications identified in this work – spectral editing and SAR reduction.

References

[1] P. O. Boykin, T. Mor, V. Roychowdhury, F. Vatan, and R. Vrijen. "Algorithmic Cooling and Scalable NMR Quantum Computers". In: *Proc. Natl. Acad. Sci. USA* 99.6 (2002), pp. 3388–3393.

[2] J. M. Fernandez, S. Lloyd, T. Mor, and V. Roychowdhury. "Algorithmic Cooling of Spins: A Practicable Method for Increasing Polarization". In: *Int. J. Quant. Inf.* 2.4 (2004), pp. 461–467.

[3] L. J. Schulman, T. Mor, and Y. Weinstein. "Physical limits of heat-bath algorithmic cooling". In: *Phys. Rev. Lett.* 94 (2005), p. 120501.

[4] L. J. Schulman, T. Mor, and Y. Weinstein. "Physical limits of heat-bath algorithmic cooling". In: *SIAM J. Comp.* 36 (2007), pp. 1729–1747.

[7] CHESS applies, in vivo, a selective 90° pulse on the target spin (commonly the large water peak at 4.7 ppm), followed by three simultaneous dephasing gradients (along each axis), pulsed gradients that spatially vary the phase of transverse magnetization. CHESS was implemented in Ref [56] by a hyperbolic secant of narrow bandwidth (~ 1.3 kHz) centered at 3.9 ppm.

[8] Taking these factors into account, AC could still yield significant cooling, particularly in solid-state NMR.

[9] A non-intuitive solution could be obtained by numerical optimization

[5] B. Criger, O. Moussa, and R. Laflamme. "Study of Multiple Rounds of Error Correction in Solid State NMR QIP". In: (2011). eprint: arxiv:quant-ph/1103.4396.

[6] A. Blank. "Scheme for a spin-based quantum computer employing induction detection and imaging". In: *Quantum Inf. Processing* 12 (2013), pp. 2993–3006.

[7] J. Baugh, J. Chamilliard, N. M. Chandrashekar, M. Ditty, A. Hubbard, R. Laflamme, M. Laforest, D. Maslov, O. Moussa, C. Negrevergne, M. P. da Silva, S. Simmons, C. A. Ryan, D. G. Cory, J. S. Hodges, and C. Ramanathan. "Quantum information processing using nuclear and electron magnetic resonance: review and prospects". In: (2007). eprint: arxiv:quant-ph/0710.1447.

[8] G. Brassard, Y. Elias, J. M. Fernandez, H. Gilboa, J. A. Jones, T. Mor, Y. Weinstein, and L. Xiao. "Experimental Heat-Bath Cooling of Spins". In: (2005). eprint: arxiv:quant-ph/0511156v1.

[9] T. Mor, V. Roychowdhury, S. Lloyd, J. M. Fernandez, and Y. Weinstein. *US patent No. 6,873,154*. 2005.

[10] Y. Elias, H. Gilboa, T. Mor, and Y. Weinstein. "Heat-bath cooling of spins in two amino acids". In: *Chem. Phys. Lett.* 517 (2011), pp. 126–131.

[11] Y. Atia, Y. Elias, T. Mor, and Y. Weinstein. "Quantum Computing Gates via Optimal Control". Submitted to Int. J. Quant. Inf.

[12] Y. Atia, Y. Elias, T. Mor, and Y. Weinstein. "Experimental Algorithmic Cooling in Liquid State NMR". manuscript in preparation.

[13] O. W. Sørensen. "Polarization transfer experiments in high-resolution NMR spectroscopy". In: *Prog. Nucl. Mag. Res. Spec.* 21 (1989), pp. 503–569.

[14] L. J. Schulman and U. V. Vazirani. "Scalable NMR Quantum Computation". In: *ACM Symposium on the Theory of Computing (STOC): Proceedings*. 1999, pp. 322–329.

[15] Y. Elias, J. M. Fernandez, T. Mor, and Y. Weinstein. "Algorithmic Cooling of Spins". In: *Isr. J. Chem.* 46 (2006), pp. 371–391.

[16] Y. Elias, J. M. Fernandez, T. Mor, and Y. Weinstein. "Optimal algorithmic cooling of spins". In: *Lecture notes in computer science: Proceedings of Unconventional Computation*. Vol. 4618. Springer publishing, 2007, pp. 2–26.

[17] F Rempp, M Michel, and G Mahler. "A cyclic cooling algorithm". In: *Phys. Rev. A* 76 (2007), p. 032325.

[18] M. J. Henrich, F. Rempp, and G Mahler. "Quantum thermodynamic Otto machines: A spin-system approach". In: *Eur. Phys. J. Spec. Top.* 151 (2007), pp. 157–165.

[19] H Weimer, M. J. Henrich, F Rempp, H Schröder, and G Mahler. "Local effective dynamics of quantum systems: A generalized approach to work and heat". In: *Europhys. Lett.* 83.3 (2008), p. 30008.

[20] H. J Briegel and S. Popescu. "Entanglement and intra-molecular cooling in biological systems? - A quantum thermodynamic perspective". In: (2008). eprint: arXiv:quant-ph/0806.4552v1.

[21] N. Linden, S. Popescu, and P Skrzypczyk. "How small can thermal machines be? The smallest possible refrigerator". In: *Phys. Rev. Lett.* 105 (2010), p. 130401.

[22] O. C. O. Dahlsten, R. Renner, E. Rieper, and V. Vedral. "Inadequacy of von Neumann entropy for characterizing extractable work". In: *New J. Phys.* 13 (2011), p. 053015.

[23] R. Renner. "Thermodynamics: The fridge gate". In: *Nature* 482 (2012), pp. 164–165.

[24] W. S. Bakr, P. M. Preiss, M. E. Tai, R. Ma, J. Simon, and M. Greiner. "Orbital excitation blockade and algorithmic cooling in quantum gases". In: *Nature* 480 (2011), pp. 500–503.

[25] S. Simmons, R. M. Brown, N. V. A. Helge Riemann, P. Becker, H.-J. Pohl, M. L. W. Thewalt, K. M. Itoh, and J. J. L. Morton. "Entanglement in a solid-state spin ensemble". In: *Nature* 470 (2011), pp. 69–72.

[26] S. Lloyd. "Quantum optics: Cool computation, hot bits". In: *Nat. Photonics* 8 (2014), pp. 90–91.

[27] J.-S. Xu, M.-H. Yung, X.-Y. Xu, S. Boixo, Z.-W. Zhou, C.-F. Li, A. Aspuru-Guzik, and G.-C. Guo. "Demon-like algorithmic quantum cooling and its realization with quantum optics". In: *Nat. Photonics* 8 (2014), pp. 113–118.

[28] Y. Elias, T. Mor, and Y. Weinstein. "Semioptimal practicable algorithmic cooling". In: *Phys. Rev. A* 83 (2011), p. 042340.

[29] J. Baugh, O. Moussa, C. A. Ryan, A. Nayak, and R. Laflamme. "Experimental implementation of heat-bath algorithmic cooling using solid-state nuclear magnetic resonance". In: *Nature* 438 (2005), pp. 470–473.

[30] C. A. Ryan, O. Moussa, J. Baugh, and R. Laflamme. "Spin based heat engine: demonstration of multiple rounds of algorithmic cooling". In: *Phys. Rev. Lett.* 100 (2008), p. 140501.

[31] Y. Atia. "Algorithmic Cooling of Spins by Optimal Control". MA thesis. Comp. Sci. Dept., Technion - Israel Institute of Technology, 2013.

[32] P. Kaye. "Cooling Algorithms Based on the 3-bit Majority". In: *Quantum Information Processing* 6 (2007), pp. 215–322.

[33] C. E. Mountford, P. Stanwell, A. Lin, S. Ramadan, and B. Ross. "Neurospectroscopy: The Past, Present and Future". In: *Chem. Rev.* 110 (2010), p. 3060.

[34] M. van der Graaf. "In vivo magnetic resonance spectroscopy: basic methodology and clinical applications". In: *Eur. Biophys. J.* 39 (2010), p. 527.

[35] B. Ross, A. Lin, K. Harris, P. Bhattacharya, and B. Schweinsburg. "Clinical experience with 13C MRS in vivo". In: *NMR Biomed.* 16 (2003), p. 358.

[36] R. Gruetter, G Adriany, I.-Y. Choi, P.-G. Henry, H. Lei, G. Öz, B. Ross, A. Lin, K. Harris, P. Bhattacharya, and B. Schweinsburg. "Localized in vivo 13C NMR spectroscopy of the brain". In: *NMR Biomed.* 16 (2003), p. 313.

[37] T. B. Rodrigues and S. Cerdán. "13C MRS: An outstanding tool for metabolic studies". In: *Concepts Magn. Reson. Part A* 27A (2005), p. 1.

[38] R. A. de Graaf, D. L. Rothman, and K. L. Behar. "State of the art direct ^{13}C and indirect ^{1}H-[^{13}C] NMR spectroscopy *in vivo*. A practical guide". In: *NMR Biomed.* 24 (2011), p. 958.

[39] D. L. Rothman, H. M. De Feyter, R. A. de Graaf, G. F. Mason, and K. L. Behar. "^{13}C MRS studies of neuroenergetics and neurotransmitter cycling in humans". In: *NMR Biomed.* 24 (2011), p. 943.

[40] S. Xu and J. Shen. "Studying enzymes by in vivo ^{13}C magnetic resonance spectroscopy". In: *Prog. Nucl. Mag. Res. Spec.* 55 (2009), p. 266.

[41] N. Sailasuta, T. T. Tran, K. C. Harris, and B. D. Ross. "Swift Acetate Glial Assay (SAGA): An accelerated human 13C MRS brain exam for clinical diagnostic use". In: *J. Mag. Reson.* 207 (2010), p. 352.

[42] N. Sailasuta, K. C. Harris, T. T. Tran, and B. D. Ross. "Minimally invasive biomarker confirms glial activation present in Alzheimer's disease: a preliminary study". In: *Neuropsychiatric Disease and Treatment* 7 (2011), p. 495.

[43] J. Kurhanewicz, D. B. Vigneron, K. Brindle, E. Y. Chekmenev, A. Comment, C. H. Cunningham, R. J. DeBerardinis, G. G. Green, M. O. Leach, S. R. Sunder, R. R. Rizi, B. D. Ross, W. S. Warren, and C. R. Malloy. "Analysis of cancer metabolism by imaging hyperpolarized nuclei: prospects for translation to clinical research". In: *Neoplasia* 13 (2011), p. 81.

[44] S. J. Nelson, J. Kurhanewicz, D. B. Vigneron, P. E. Z. Larson, A. L. Harzstark, M. Ferrone, M. van Criekinge, J. W. Chang, R. Bok, I. Park, G. Reed, L. Carvajal, E. J. Small, P. Munster, V. K. Weinberg, J. H. Ardenkjaer-Larsen, A. P. Chen, R. E. Hurd, L.-I. Odegardstuen, F. J. Robb, J. Tropp, and J. A. Murray. "Metabolic Imaging of Patients with Prostate Cancer Using Hyperpolarized [1-^{1}3C]Pyruvate". In: *Sci. transl. med.* 5 (), 198ra108.

[45] B. Ross, P. Bhattacharya, S. Wagner, T. Tran, and N. Sailasuta. "Hyperpolarized MR Imaging: Neurologic Applications of Hyperpolarized Metabolism". In: *Am. J. Neuroradiol.* 31 (2010), p. 24.

[46] Y. Xiang and J. Shen. "Simultaneous detection of cerebral metabolism of different substrates by *in vivo* ^{13}C MRS". In: *J. Neurosci. Met.* 198 (2011), p. 8.

[47] F. Boumezbeur, G. F. Mason, R. A. de Graaf, K. L. Behar, G. W. Cline, G. I. Shulman, D. L. Rothman, K. F. Petersen, and M. van der Graaf. "Altered brain mitochondrial metabolism in healthy aging as assessed by in vivo magnetic resonance spectroscopy". In: *J. Cereb. Blood Flow & Metab.* 30 (2010), p. 211.

[48] S. Li, Y. Zhang, S. Wang, J. Yang, A. M. Ferraris, A. Farris, C. Johnson, S. Fox, R. Innis, and J. Shen. "In vivo ^{13}C magnetic resonance spectroscopy of human brain on a clinical 3T scanner using [2-^{13}C]glucose infusion and low-power stochastic decoupling". In: *Magn. Reson. Med.* 62 (2009), p. 565.

[49] D. K. Deelchand, K. Uğrbil, and P.-G. Henry. "Investigating Brain Metabolism at High Fields using Localized 13C NMR Spectroscopy without 1H Decoupling". In: *Magn. Reson. Med.* 55 (2006), p. 279.

[50] S. Li, J. Yang, and J. Shen. "Novel strategy for cerebral 13C MRS using very low RF power for proton decoupling". In: *Magn. Reson. Med.* 57 (2007), p. 265.

[51] N. Sailasuta, L. W. Robertson, K. C. Harris, A. L Gropman, P. S. Allen, and B. D. Ross. "Clinical NOE 13C MRS for neuropsychiatric disorders of the frontal lobe". In: *J. Mag. Reson.* 195 (2008), p. 219.

[52] Y. Xiang and J. Shen. "Windowed Stochastic Proton Decoupling for In Vivo ^{13}C Magnetic Resonance Spectroscopy with Reduced RF Power Deposition". In: *J. Mag. Reson. Imaging* 34 (2011), p. 968.

[53] T. W.-M. Fan and A. N. Lane. "Structure-based profiling of metabolites and isotopomers by NMR". In: *Prog. Nucl. Mag. Res. Spec.* 52 (2008), p. 69.

[54] W. Wiechert, M. Möllney, N. Isermann, and A. A. Wurzel M. an de Graaf. "Bidirectional Reaction Steps in Metabolic Networks: III. Explicit Solution and Analysis of Isotopomer Labeling Systems". In: *Biotechnology and Bioengineering* 66 (1999), p. 69.

[55] D. P. Burum and R. R. Ernst. “Net polarization transfer via a J-ordered state for signal enhancement of low-sensitivity”. In: *J. Mag. Reson.* 39 (1980), p. 163.

[56] J. Shen, J. Yang, I.-Y. Choi, S. S. Li, and Z. Chen. “A new strategy for in vivo spectral editing. Application to GABA editing using selective homonuclear polarization transfer spectroscopy”. In: *J. Mag. Reson.* 170 (2004), p. 290.

[57] S. Xu and J. Shen. “In vivo dynamic turnover of cerebral 13C isotopomers from [U-13C]glucose”. In: *J. Mag. Reson.* 182 (2006), p. 221.

[58] O. W. Sørensen. “The entropy bound as a limiting case of the universal bound on spin synamics. Polarization transfer in $I_N S_M$ spin systems”. In: *J. Mag. Reson.* 93 (1991), p. 648.